

Report on the Acceptance Test of the CRI Y-MP 8128 10 February - 12 March 1990

RND 90-002

*Russell Carter*¹

Computer Sciences Corporation
NASA Ames Research Center
Moffett Field, CA 94035, USA

Abstract

The NAS Numerical Aerodynamic Simulation Facility's HSP 2 computer system, a CRI Y-MP 832 SN #1002, underwent a major hardware upgrade in February of 1990. The 32 MWord, 6.3 ns mainframe component of the system was replaced with a 128 MWord, 6.0 ns CRI Y-MP 8128 mainframe, SN #1030. As per NASA contract NAS2-12762, a 30 day Acceptance Test of the computer system was performed by the NAS RND HSP group from 08:00 February 10, 1990 to 08:00 March 12, 1990. Overall responsibility for the RND HSP Acceptance Test was assumed by Duane Carbon. The terms of the contract required that the SN #1030 achieve an effectiveness level of greater than or equal to ninety (90) percent for 30 consecutive days within a 60 day time frame. After the first thirty days, the effectiveness level of SN #1030 was 94.4 percent, hence the acceptance test was passed.

A. Effectiveness Level Determination

As defined in contract NAS2-12762, the effectiveness level of the system is computed by dividing the operational-use time by the sum of the operational-use time plus the system failure downtime. Operational use time (OPUSE) is the actual time that all processors are available to perform the actual or simulated Government workload. Evidence of whether or not the #1030 was correctly processing the Government workload was provided in part by RND HSP test codes, described in section B. System failure downtime (SYSFAIL) is the time in which the system is unusable to process the Government workload at the required performance levels due to Contractor-supported equipment or standard software failure. System downtime due to normal Preventative Maintenance (PM), dedicated time, or NAS Operations Branch (RNS) errors is not counted to either OPUSE or SYSFAIL, but instead counted to null time (NULL).

¹ This work was supported by NASA Contract No. NAS2-12961 while the author was an employee of Computer Sciences Corporation under contract to the Numerical Aerodynamic Simulation Systems Division at NASA Ames Research Center.

Effectiveness level (EFF) is then expressed as the following formula:

$$EFF = \frac{OPUSE}{(OPUSE + SYSFAIL)}$$

Note that the total hours used to compute the effectiveness level (OPUSE + SYSFAIL) is less than or equal to the total hours in 30 consecutive days (OPUSE + SYSFAIL + NULL).

During the acceptance test, the effectiveness level of the SN #1030 was computed on a daily basis. The number of hours of OPUSE, SYSFAIL, and NULL were determined by the examination of the wilbur:/usr/unsupported/bin/stat (stat) utility, the RNS daily Operations Log located in the operators' room, and the ymp:/etc/special.log (special.log).

The determination of downtime for most cases required examination of the Operations Log for each twenty-four hour period. Problems with the Y-MP recorded in the Operations Log were noted and downtime charged to SYSFAIL was counted from the time that the CRI Field Engineer (FE) was notified. However, in one case, downtime was assessed from the time the system failed to correctly process the Government workload (see Appendix B, (26)).

The *stat* utility and *special.log* were used to determine the start of uptime (OPUSE). The Operations Log record of system boot time was used as a reference when the *special.log* was inspected. If any special queue jobs ran, uptime was counted from the start of the first special queue job. Two entries from /etc/special.log follow:

Sample of /etc/special.log:

```
Wed Feb 21 19:02:23 PST 1990
Starting special job /etc/special/runwf for storaasl.g1324 at Wed Feb 21
19:02:24 PST 1990
Finished special job /etc/special/runwf for storaasl.g1324 at Wed Feb 21
19:17:24 PST 1990
```

```
Tue Feb 27 01:10:19 PST 1990
Starting special job /etc/special/runwf for storaasl.g1324 at Tue Feb 27
01:10:20 PST 1990
Finished special job /etc/special/runwf for storaasl.g1324 at Tue Feb 27
01:25:20 PST 1990
```

If no special queue jobs ran, uptime was determined by inspection of output from the *stat* utility. A sample of *stat* output is provided on the following page.

Sample output from *stat*:

	NAME	SEQ	DATE	TIME	%IDLE	MEM	JOBS	DAEMON	USERS	PROCS	LOAD	ERR
	reynolds	1065	03/08	19:33:28	323.33	100	11	1	23	203	--	---
	reynolds	1066	03/08	19:34:28	318.33	44	9	1	23	186	--	---
	reynolds	1067	03/08	19:35:28	50.00	90	2	1	22	154	--	---
	reynolds	1068	03/08	19:36:29	80.32	60	0	0	23	143	--	---
	reynolds	1	03/08	22:20:10	??	6	0	0	0	45	--	---
	reynolds	2	03/08	22:21:10	433.33	106	8	1	0	106	--	---
	reynolds	3	03/08	22:22:11	221.66	107	8	1	2	86	--	---
	reynolds	4	03/08	22:23:11	305.00	104	8	1	1	84	--	---
	reynolds	5	03/08	22:24:11	393.33	106	8	1	2	87	--	---

During downtime, the *stat* utility reports a '0' in the DAEMON column, indicating the NQS daemon is not up. Uptime was counted from the time that *stat* reported a '1' in the DAEMON column.

Daily and Cumulative percent OPUSE and SYSFAIL were daily presented in tables and graphs and distributed to members of NAS RND and CRI. Time credited to SYSFAIL was considered tentative until approved by the Contracting Officer's Technical Representative (COTR), John Barton. The COTR's decision was made using input from daily meetings between Government and CRI representatives. The daily log for the final day of the Acceptance Test, March 12, 1990, follows as Appendix A. Details for each numbered downtime are provided in Appendix B.

B. RND HSP Test Codes

RND HSP Group members Robert Bergeron, Russell Carter, Robert Ciotti, Teresa Griffie, Eugene Miya, and Douglas Pase, provided codes to test hardware and software functions of the SN #1030. System hardware components explicitly tested included memory and CPU integrity, I/O, SSD and memory swapping functionality, and the multitasking facilities (semaphores). Software functionality was tested by exercising the C and FORTRAN compilers, the multitasking libraries (autotasking, microtasking and macrotasking), and UNICOS system calls. RND HSP group members ran their codes at frequent and periodic intervals, noting all failures. Failures were investigated, and if attributable to a failure in the SN #1030, were reported at the daily CRI/RND HSP meetings.

The test codes uncovered a number of problems with the SN #1030. There was an initial problem configuring swap. IOS striping did not allow the configuration desired by NAS, which led CRI to configure the system with minimal swap, which was later increased but without striping. Mainframe striping was later added and led to swapping problems found later with large codes. The solution implemented by CRI consisted of increasing the amount of swapspace and the addition of a 'bigproc' parameter to the schedv utility. The latter change prevents codes with memory requirements larger than bigproc from swapping. Intermittent failures of several large memory codes run by R. Bergeron and E. Miya led to the discovery by CRI of a bug in the implementation of mainframe striping of swap. CRI eliminated the problem with a kernel fix.

D. Pase noticed that multitasked codes experience increasingly large performance


degradations as the number of virtual processors is increased from 30 to 40 and up.

R. Ciotti found several problems with the SN #1030. He discovered a problem with checkpointed jobs and the *timf* routine. He found that jobs restarted after a shutdown dumped core when subsequent calls to *timf* were made. CRI applied fixes to prevent core dumps, but final resolution awaits UNICOS 6.0. In addition, R. Ciotti notified CRI that *acctcom* needed recompilation with the new system parameters, and the initial kernel memory usage of 8.2 MWords was excessive. CRI reduced kernel size to 3.7 MWords. R. Ciotti also noted that the new scheduler (with *bigproc*) exhibits poor behavior when swap is overloaded.

The final incident is an archetypal example of human error and the need for validation of computer systems. D. Pase twice observed, on 8 March 90, single bit errors in the output of one of his codes. CRI noted that this was impossible with a correctly working system. Subsequent investigation by CRI revealed that the SN #1030 had run the Government workload since the last PM with memory error checking (SECDED) accidentally turned off. Since the validity of computations performed during the time SECDED was off was impossible to determine, SN #1030 was declared down for the entire period.

A summary of RND HSP group test code runs is provided in Appendix C. Brief code descriptions are provided in Appendix D. A summary of the cpu time used by the codes appears in Appendix E.

Cumulative YMP Percent UP/DOWN Time
March 12, 1990



Cumulative YMP Percent UP/DOWN Time
March 12, 1990

